arXiv:1405.5139v2 [math.LO] 24 May 2014

Algorithmic identification of probabilities is hard

Laurent Bienvenu[1], Benoît Monin[2], and Alexander Shen[3]

[1] Laboratoire Poncelet, laurent.bienvenu@computability.fr
[2] LIAFA benoit.monin@liafa.univ-paris-diderot.fr
[3] LIRMM alexander.shen@lirmm.fr

Abstract. Suppose that we are given an infinite binary sequence which is random for a Bernoulli measure of parameter p. By the law of large numbers, the frequency of zeros in the sequence tends to p, and thus we can get better and better approximations of p as we read the sequence. We study in this paper a similar question, but from the viewpoint of inductive inference. We suppose now that p is a computable real, but one asks for more: as we are reading more and more bits of our random sequence, we have to eventually guess the exact parameter p (in the form of a Turing code). Can one do such a thing uniformly on all sequences that are random for computable Bernoulli measures, or even on a 'large enough' fraction of them? In this paper, we give a negative answer to this question. In fact, we prove a very general negative result which extends far beyond the class of Bernoulli measures.

1 Introduction

1.1 Learnability of sequences

The study of learnability of recursive sets is concerned with the following problem: suppose that we have a sequence of bits $X = X(0)X(1)X(2), ...$ which is computable. That is, there is a total Turing machine which can write on an infinite output tape the sequence X. (Here *total* means that the machine does produce an infinite sequence; not all machines have this property). Suppose also that we are given access to a finite prefix of X

$$X \upharpoonright n = X(0) \dots X(n-1)$$

Our goal is to guess the code[4] of a Turing machine which computes X, and we want the guessing process to be itself computable. One way to visualize this problem is to imagine we have a Turing machine M in a blackbox, and can only observe its output. The goal is to guess what machine is in the blackbox. One needs to keep in mind that for each machine M, there are infinitely many other machines that are extensionally equivalent (produce the same output), so it is impossible to guess the exact machine M in the blackbox, but guessing an

[4] Turing machines are finite objects and can be computably listed. If one fixes such an enumeration $M_0, M_1, ...$ once and for all, we say that the *code* of M_e is e

equivalent machine M' would count as success: such a machine would indeed be a good *explanation* for the observed sequence. For example, after having observed the finite string

$$11001001000011111011010101000$$

we might want to guess that the machine M writes the binary expansion of π and thus propose the code of a machine which does exactly that.

Of course, finitely many data points are probably not enough to guess the entire sequence (in the above example, the observed finite string does not rule out the possibility that the sequence is the binary expansion of $47627751/15160384$). The hope is that, as we gain access to more and more data points, we will *eventually* figure out what the sequence X is, and how it was generated. More precisely, we hope to have a computable function $\mathfrak{A}$ such that for every X computable, the sequence

$$\mathfrak{A}(X\upharpoonright 1),\ \mathfrak{A}(X\upharpoonright 2),\ \mathfrak{A}(X\upharpoonright 3),...$$

eventually produces a code for a Turing machine computing X, which is what is referred to as *identification in the limit*. How quickly $\mathfrak{A}$ figures out a good code will depend on the sequence X. Note that there are two possible levels of success our procedure can have:

- Strong success: for every computable X, the above sequence converges to a single value e, and e is the code of a machine which produces X.
- Weak success: for every computable X, all but finitely many terms of the above sequence are codes of machines which produce X; however, these codes may be different.

The first type of success is often referred to as *exact (EX)* and the second type as *behaviorally correct (BC)*. Either way, such an algorithm $\mathfrak{A}$ does not exist in general, the main obstacle being that certain Turing machines are not total, and distinguishing total machines from non-total machines cannot be done computably.

Thus, any algorithm will fail to identify certain total computable function and a natural question to ask is what classes of total computable functions can be identified by a single algorithm $\mathfrak{A}$. For example, there *is* an algorithm $\mathfrak{A}$ which identifies correctly (even in the EX sense) any X in

$$\mathcal{U}_0 = \{X \text{ computable such that } \forall^\infty n\ X(n) = 0\}$$

We refer the reader to [ZZ08] for a detailed survey of learnability of recursvie functions.

1.2 Learnability of probability measures

Recently, Chater and Vitanyi [CV13] proposed to study a related problem: Suppose now that instead of a total deterministic Turing machine, the blackbox contains an *almost total probabilistic* Turing machine M. By "almost total" Turing

machine, we mean a Turing machine with probabilistic transitions, which almost surely produces an infinite binary sequence X. The machine being probabilistic, X is now a random variable and every almost total probabilistic machine induces a probability measure μ_M over the space 2^ω of infinite binary sequences. Again, our goal is to – in the limit – make a good guess as to what machine is in the box, i.e, give a reasonable explanation for the observed sequence X. For example, if we observe the sequence

$$0001111111100001100000000011111111111111$$

a reasonable guess would be to say that M is a probabilistic machine which produces sequence of bits where the $(n+1)$-th term is equal to the n-th term with probability $4/5$ and equal to the n-th term plus one with probability $1/5$.

In this context, what should count as a good guess for an observed sequence? Like for deterministic machines, we cannot be required to guess the exact machine M, and should also be allowed to guess the code of another machine M' which behaves exactly like M (has the same probability distribution over outputs: $\mu_M = \mu_{M'}$). But even this is too much to ask for. Indeed, suppose X is the observed sequence and e is the code of a probabilistic machine M which we would like to count as 'good guess' for X. Consider the probabilistic machine M' which does the following: it first makes a random $0/1$ choice with probability $(1/2, 1/2)$. In the first case, it outputs the sequence $0, 0, 0, ...$ (making no further probabilistic choice). In the second case, it starts behaving like M. Since the machine M' behaves like M with probability $1/2$, and we only have a single observed sequence X, then if X is a plausible output of M, it certainly is a plausible output for M', even though M and M' have different behaviours.

This seems to call into question whether one can give a satisfactory formalization of 'good guess'. Fortunately, it can indeed be done, thanks to the theory of algorithmic randomness. As Chater and Vitanyi recall, given an almost total probabilistic Turing machine with output distribution μ, there is a widely accepted definition of what are the typical outputs the machine are, namely, the sequences which are Martin-Löf random with respect to the measure μ. These are the sequences which pass all effective statistical tests for the measure μ, also known as μ-Martin-Löf tests.

Thus our goal is to give an algorithm $\mathfrak{A}$ such that for every almost total probabilistic machine M with output distribution μ_M, for all μ_M-almost all X, the sequence $\mathfrak{A}(X \upharpoonright 1), \mathfrak{A}(X \upharpoonright 2), \mathfrak{A}(X \upharpoonright 3), ...$ identifies in the limit an almost total probabilistic machine M' such that X is $\mu_{M'}$-Martin-Löf random. Chater and Vitanyi suggest that this can be achieved in the strongest sense (EX): the $\mathfrak{A}(X \upharpoonright n)$ converge to a single code for a machine M'). The main result of this paper is that it cannot in fact be achieved, even if we require only that:

- The algorithm succeeds in a very weak sense, i.e., that there exists a $c > 0$, such that all but finitely n, every $\mathfrak{A}(X \upharpoonright n)$ is a code of a machines M' such that X is $\mu_{M'}$-Martin-Löf random with randomness deficiency $< c$.
- For every machine M the algorithm succeeds with μ_M-probability at least $\delta > 0$ (for some δ independent of M), instead of probability 1.

– The algorithm succeeds only on inputs that are μ_M-random for a restricted class of machines M, as long as this class satisfies reasonable assumptions (see next section for the precise formulation).

The proof of this result is quite involved. The rest of the paper is devoted to presenting this proof, and discuss further consequences. We assume that the reader is familiar with algorithmic randomness and Kolmogorov complexity. The most useful references for our purposes are [Gác05] and [LV08].

2 Identifying measures

2.1 Background and notation

Let us start by providing some notation and background.

We denote by 2^ω the set of infinite binary sequence and by $2^{<\omega}$ the set of finite binary sequences (or *strings*). The length of an element σ of $2^{<\omega}$ is denoted by $|\sigma|$. If X is either an element of 2^ω or an element of $2^{<\omega}$ of length at least n, then $X(n-1)$ denotes the n-th bit of X and $X \upharpoonright n$ denotes the finite sequence $\sigma = X(0)X(1)\dots X(n-1)$. If σ is a member of $2^{<\omega}$ and X is either an infinite sequence or sequence of length at least $|\sigma|$, we say that σ is a *prefix* of X, which we denote $\sigma \preceq X$, if $X \upharpoonright |\sigma| = \sigma$.

The space 2^ω is endowed with the distance d defined by

$$d(X,Y) = 2^{-\min\{n : X(n) \neq Y(n)\}}$$

This distance is compatible with the product topology, which is the topology generated by *cylinders*, where a cylinder is a set of type

$$[\sigma] = \{X \in 2^\omega \ : \ \sigma \preceq X\}$$

A cylinder is both open and closed (= clopen). Thus, any finite union of cylinders is also clopen. It is easy to see, by compactness, that the converse holds: the only clopen subsets of 2^ω are the finite union of cylinders. We say that a clopen set C has *granularity at most* n if it can be written as a finite union of cylinders $[\sigma]$ with all σ's of length $\leq n$. We denote by Γ_n the set of cylinders of granularity at most n.

The space of Borel probability measures over 2^ω is denoted by $\mathcal{M}(2^\omega)$. It is equipped with the weak topology. Several classical distances are compatible with this topology; for our purposes, it will be convenient to use the distance ρ, constructed as follows: For $\mu, \nu \in 2^\omega$ for all n, let $\rho_n(\mu,\nu)$ be the quantity

$$\rho_n(\mu,\nu) = \max_{C \in \Gamma_n} |\mu(C) - \nu(C)|$$

and then set

$$\rho(\mu,\nu) = \sum_n 2^{-n} \rho_n(\mu,\nu)$$

The open (resp. closed) ball $\mathcal{B}$ of center μ and radius r is the set of measures ν such that $\rho(\mu,\nu) < r$ (resp. $\rho(\mu,\nu) \leq r$). Note that for any ν in this open (resp. closed) ball, if C is a clopen set of granularity at most n, then $|\mu(C)-\nu(C)| < 2^n\, r$ (resp. $\leq 2^n\, r$). The distance ρ makes $\mathcal{M}(2^\omega)$ a computable compact metric space, th ecomputable points of which are called computable probability measures and a measure is computable if and only if it is the output distribution of some almost total probabilistic Turing machine [Gác05]. Being a computable metric space, one can define partial computable functions from some discrete space $\mathcal{X}$ (such as $\mathbb{N}$) to $\mathcal{M}(2^\omega)$ via type-2 computability: a partial function $f :\subseteq \mathcal{X} \to \mathcal{M}(2^\omega)$ is partial computable if there is a computable operator g which on every input $x \in \mathcal{X}$ enumerates a list of balls $\mathcal{B}_1$, $\mathcal{B}_2$, ... (which may be finite) of $\mathcal{M}(2^\omega)$ such that $\mathcal{B}_{i+1} \subseteq \mathcal{B}_i$, the diameter of $\mathcal{B}_i$ is less than 2^{-i} and such that for every x in the domain of f, the list of enumerated balls is infinite and their intersection is the singleton $\{f(x)\}$. Let us introduce two non-standard, but important in this paper, pieces of terminology: having fixed the operator g associated to f, we write $\mathrm{err}(f(x)) < \varepsilon$ to mean that the list of balls produced by g on input x contains a ball of radius less than ε (the justification for this notation is that when such a ball is enumerated, should $f(x)$ be defined, we know its value with error at most ε for the distance ρ). When the operator g on input x enumerates an empty list of balls, we say that g is *null* on input x.

We denote by K the prefix-free Kolmogorov complexity function. Given a computable measure μ, we call *randomness deficiency of X with respect to μ* the quantity

$$\mathbf{d}(X|\mu) = \sup_n \left[\log \frac{1}{\mu([X\restriction n])} - \mathrm{K}(X\restriction n)\right]$$

$X \in 2^\omega$ is μ-Martin-Löf random (or μ-random for short) if $\mathbf{d}(X|\mu) < \infty$. This definition is slightly non-standard: one usually uses $\mathrm{K}(X\restriction n|\mu)$ instead of $\mathrm{K}(X\restriction n)$ in the above expression, see [LV08]. However, the above is enough for our purposes and will actually yield a slightly stronger result. We say that two measures μ and ν are *orthogonal* if there is a set of measure 1 for μ and measure 0 for ν.

If $\mathcal{B}$ is a ball (open or closed) of $\mathcal{M}(2^\omega)$, of center μ and radius r, we define the *estimated deficiency* of X relative to $\mathcal{B}$ by

$$\mathbf{ed}(X|\mathcal{B}) = \sup_n \left[\log \frac{1}{\mu([X\restriction n]) + 2^n\, r} - \mathrm{K}(X\restriction n)\right]$$

Note that $\mathbf{d}(X|\nu) \geq \mathbf{ed}(X|\mathcal{B})$ for all $\nu \in \mathcal{B}$ and for a fixed pair (X,μ), that $\lim_{\mathcal{B}\to\mu} \mathbf{ed}(X|\mathcal{B}) = \mathbf{d}(X|\mu)$.

The next lemma will be useful in the sequel.

Lemma 1 (Randomness deficiency lemma). *Let $\mathcal{B} \subseteq \mathcal{M}_0$ be a ball of center μ (computable measure) and rational radius $\leq r$, and C a clopen set of granularity at most n. Then for all $X \in C$:*

$$\mathbf{ed}(X|\mathcal{B}) \geq -\log(\mu(C) + 2^n r) - K(C,\mu,r,n) - O(1)$$

Proof. Let $m = -\log(\mu(C) + 2^n r)$. Let $\{\sigma_i\}_{1 \leq i \leq 2^n}$ be a set of pairwise disjoint strings of length less than n which induces C. Then let us create the bounded request set (a.k.a. Kraft-Chaitin set, see [Nie09,DH10]):

$$\left\{ \sigma_i \ , \ \log \frac{1}{\mu(\sigma_i) + 2^n r} - m \right\}_{i \in \mathbb{N}}$$

Then we have that

$$\begin{aligned} \textstyle\sum_i 2^{-log \frac{1}{\mu(\sigma_i)+2^n r} + m} &= 2^m (\textstyle\sum_i \mu(\sigma_i) + 2^n r) \\ &= 2^m (\mu(C) + 2^n r) \\ &\leq 1 \end{aligned}$$

The bounded request set can be uniformly computed from C, μ, r, n. Thus for every i,

$$\mathrm{K}(\sigma_i | C, \mu, r, n) \leq \log \frac{1}{\mu(\sigma_i) + 2^n r} - m + O(1)$$

and thus

$$\mathrm{K}(\sigma_i) \leq \log \frac{1}{\mu(\sigma_i) + 2^n r} - m + K(C, \mu, r, n) + O(1)$$

The result follows from the definition of $\mathbf{ed}(X|\mathcal{B})$, using the fact that every $X \in C$ has some σ_i as a prefix. □

2.2 The main theorem

Now that we have set the stage, we can formalize the idea of inductive inference of measures. To get the strongest possible theorem, we need to work in the setting of behaviourally correct learning, that is when the outputs of the algorithm do not have to converge to a single code. In this setting, outputting a code of a computable measure is equivalent to computing the measure in the type-2 sense, so we will call a (measure)-learning algorithm a partial computable function $\mathfrak{A} :\subseteq 2^{<\omega} \to \mathcal{M}(2^\omega)$. We say that $\mathfrak{A}$ *BC-succeeds* on a sequence $X \in 2^\omega$ if for almost every n, $\mathfrak{A}(X \upharpoonright n) = \mu$, for some computable measure μ with respect to which X is Martin-Löf random. Our main result, in its weak form, is the following.

Theorem 2. *There is no algorithm $\mathfrak{A}$ which BC-succeeds on every sequence X which is random with respect to some computable measure.*

We shall, however, prove a significantly stronger result. One could ask for example if there could be an algorithm which works on X's which are random for a specific type of measures, such as Bernoulli measures (the main particular case considered by Chater and Vitanyi). Or one could ask whether there could be an algorithm which does not necessarily work for all X's but does work for a fixed fraction of them. Finally, one could weaken the notion of desired success, not requiring that the algorithm converges to a specific measure, but asking that the algorithm gives 'good explanations' for the observed sequence from some point

on. More specifically, we say that an algorithm *BD-succeeds* (BD standing for 'bounded deficiency') on some X if for some N and almost every n, $\mathfrak{A}(X\upharpoonright n)$ is defined and is a measure with respect to which X is random with deficiency at most N. Clearly BC-success implies BD-success.

We now state our main result in its strong form.

Theorem 3. *Let $\mathcal{M}_0$ be an effectively compact subset of $\mathcal{M}(2^\omega)$ with the property that in every neighbourhood of $\mathcal{M}_0$, there are countably many pairwise effectively orthogonal computable measures. Let $\delta > 0$. There is no algorithm $\mathfrak{A}$ such that for every computable $\mu \in \mathcal{M}_0$, the μ-measure of X's on which $\mathfrak{A}$ BD-succeeds is at least δ.*

Note that the hypotheses on the class $\mathcal{M}_0$ are not very restrictive as many standard classes of probability measures have these properties. Bernoulli measures are an obvious example and thus our theorem applies to any class $\mathcal{M}_0$ in which the class of Bernoulli measures can be embedded (such as the class of Markov chains). Let us give another interesting example: for every parameter $p \in [0,1]$, consider measure μ_p associated to the stochastic process which generates a binary sequence bit-by-bit as follows: the first bit is 1, and having generated a string $\sigma 10^k$, the probability of the next bit to be 1 is equal to $p/(k+1)$. The class $\mathcal{M}(2^\omega) = \{\mu_p : p \in [0,1]\}$ satisfies the hypotheses of the theorem.

Note also that these hypotheses are not added for convenience: although they might not be optimal, they cannot be outright removed. If we do not require compactness, then the class of Bernoulli measures of *rational* parameter would qualify, but it is easy to see that this class admits an algorithm which correctly identifies each of the measures in the class with probability 1. The second condition is important too. Consider the measure ν_0 which is the measure concentrated on the sequence 0^ω (sequence consisting only of zeroes) and ν_1 the measure concentrated on 1^ω. Then the class $\mathcal{M}(2^\omega) = \{p\nu_0 + (1-p)\nu_1 \mid p \in [0,1]\}$ is indeed effectively compact, but it is obvious that there is an algorithm which BC-succeeds with probability 1 for all measures of that class (by simply outputting ν_i, where i is the first bit of the sequence being read).

3 The proof of the main theorem

The rest of the paper is devoted to proving Theorem 3. We fix a subset $\mathcal{M}_0$ of $\mathcal{M}(2^\omega)$ satisfying the hypotheses of the theorem, and also fix some $\delta > 0$ once and for all. For readability, we abbreviate BD-success by simply 'success'.

We want to show that no algorithm $\mathfrak{A}$ succeeds on a δ-fraction of X's for every $\mu \in \mathcal{M}_0$. Call such an hypothetical algorithm *δ-good*.

Let us introduce some useful notation. First, denote

$$\mathrm{SUCC}(\mathfrak{A}, N, n) = \left\{ X \in 2^\omega \ : \ (X\upharpoonright n) \in \mathrm{dom}(\mathfrak{A}) \wedge \ \mathbf{d}(X | \mathfrak{A}(X\upharpoonright n)) \leq N \right\}$$

One can see $\mathrm{SUCC}(\mathfrak{A}, N, n)$ as the set of X's on which $\mathfrak{A}$ achieves 'local success' on the prefix of length n for randomness deficiency N. Thus, the set of X's on which $\mathfrak{A}$ succeeds (which we refer to as *success set*) is $\bigcup_N \bigcap_{n \geq N} \mathrm{SUCC}(\mathfrak{A}, N, n)$.

Given our above description of type-2 computability, we identify $\mathfrak{A}$ with its 'listing operator'. Thus, we will write '$\mathcal{B} \in \mathfrak{A}(\sigma)$' to signify that on input σ, the listing operator for $\mathfrak{A}$ enumerates the ball $\mathcal{B}$ at some point. For any decreasing function $f : \mathbb{N} \to [0, 1]$ converging to 0, we define the set $\mathrm{PREC}(A, f, n)$ of points X which are 'precise' in the sense that $A(X \upharpoonright n)$ almost outputs a measure:

$$\mathrm{PREC}(\mathfrak{A}, f, n) = \{X \in 2^\omega \ : \mathrm{err}(\mathfrak{A}(X \upharpoonright n)) < f(n)\}$$

(notice that $\mathrm{PREC}(\mathfrak{A}, f, n)$ is a clopen set because the membership of X in $\mathrm{PREC}(\mathfrak{A}, f, n)$ is determined fully by the first n bits of X).

In contrast to PREC, we define the following 'nullity' sets:

$$\mathrm{NULL}(\mathfrak{A}, N, n) = \left\{X \in 2^\omega \ : \ \mathfrak{A}(X \upharpoonright n) \text{ is null}\right\}$$

and

$$\mathrm{NULL}(\mathfrak{A}, N) = \bigcap_{n \geq N} \mathrm{NULL}(\mathfrak{A}, N, n)$$

We reduce the main theorem to the following proposition.

Proposition 4 (Nullity amplification). *Let $\mathfrak{A}$ be a δ-good algorithm, N an integer, η a real and $\mathcal{B} \subseteq \mathcal{M}_0$ a non-empty ball such that $\mu(NULL(\mathfrak{A}, N)) \geq \eta$ for all $\mu \in \mathcal{B}$. Then, there is a non-empty ball $\mathcal{B}' \subseteq \mathcal{B}$, an integer $N' \geq N$ and a δ-good algorithm $\mathfrak{A}'$ such that $\mu(NULL(\mathfrak{A}', N')) \geq \eta + \delta/2$ for all $\mu \in \mathcal{B}'$.*

This proposition clearly shows that there can be no δ-good algorithm: if there was one, one could construct by induction a sequence of δ-good algorithms $\mathfrak{A}_i$, a non-increasing sequence of non-empty balls $\mathcal{B}_i$ and a non-decreasing sequence of integers N_i such that for every $\mu \in \mathcal{B}_i$, $\mu(\mathrm{NULL}(\mathfrak{A}_i, N_i)) \geq \delta + i \cdot (\delta/2)$, which is impossible for $i > 2/\delta$. Thus, all we need to do is prove this proposition.

Proof. We consider the family of effectively open sets

$$\mathcal{U}_m = \{\mu \mid (\exists n > m)\, \mu(\mathrm{PREC}(\mathfrak{A}, f, n)) > 1 - \eta - \delta/2\}$$

for $m \geq N$ and we show that only finitely many of them can be dense (as long as f tends to 0 fast) in $\mathcal{B}$. The way we do this is by proving that if $\mathcal{U}_m$ is dense for some m, then for every non-empty $\mathcal{B}' \subseteq \mathcal{B}$, one can effectively find a non-empty $\mathcal{B}'' \subseteq \mathcal{B}'$ such that for all $\mu \in \mathcal{B}''$, $\mu(\mathrm{SUCC}(\mathfrak{A}, n, n)) < \delta/2$ for some $n \geq m \geq N$. This would yield a contradiction since this would allow us to construct an effective sequence of decreasing balls $\mathcal{B}_m$, where all $\mu \in \mathcal{B}_m$ would be such that $\mu(\mathrm{SUCC}(\mathfrak{A}, n, n)) < 7\delta/8$ for some $n \geq m$, and thus the intersection

of the $\mathcal{B}_m$ would be a computable measure μ^* for which the success set of $\mathfrak{A}$ has μ^*-measure at most $7\delta/8$, a contradiction.

We first need to define the function f. The definition of $f(n)$ as a function of n will depend on a 'large enough' parameter $s = s(n)$ which we will define later as a computable function of n itself. For n fixed, let $\varepsilon = \min(2^{-n} \cdot \delta/4, r)$ where r is the radius of $\mathcal{B}$.

First, we effectively find $k(\varepsilon)$ rational balls $\mathcal{D}_1, \mathcal{D}_2, \cdots \mathcal{D}_{k(\varepsilon)}$ covering $\mathcal{M}_0$ so that for any ball of radius at least ε, one of the $\mathcal{D}_i$ is contained in this ball (To do this, enumerate all balls of rational center and radius smaller than $\varepsilon/2$. By effective compactness, one can find a finite number of them, call them $\mathcal{B}_1, \mathcal{B}_2, \cdots \mathcal{B}_{k(\varepsilon)}$ which cover the entire space. Now, let $\mathcal{A}$ be a ball of radius at least ε and μ its center. By definition, there is an i such that $\mu \in \mathcal{B}_i$, and by the triangular inequality, any member of $\mathcal{B}_i$ is a distance at most $\varepsilon/2 + \varepsilon/2 = \varepsilon$ of μ, hence $\mathcal{B}_i$ is contained in $\mathcal{A}$).

Then, for each ball $\mathcal{D}_i$, we effectively find 2^s distinct computable measures $\xi_1^{(i)}, \ldots, \xi_{2^s}^{(i)}$ and a finite family of pairwise disjoint clopen sets $V_1^{(i)}, \ldots, V_{2^s}^{(i)}$ so that $\xi_j^{(i)}(V_j^{(i)}) > 1 - \delta/8$.

To see that this can be done, observe that the conditions '$\xi_1, \ldots, \xi_s \in \mathcal{B}$', 'the V_i are disjoint', and '$\xi_i(V_i) > 1 - \varepsilon$ for all i' are all Σ_1^0 conditions. Therefore, all we need to argue is that such measures and clopen sets exist. By our assumption on $\mathcal{M}_0$, let $\xi_1, \ldots, \xi_s$ be pairwise orthogonal measures inside $\mathcal{B}$. By definition, this means that for every pair (i, j) with $i \neq j$, there exists a set $S_{i,j} \subseteq 2^\omega$ such that $\xi_i(S_{i,j}) = 1$ and $\xi_j(S_{i,j}) = 0$. For each i, let $S_i = \bigcap_{j \neq i} S_{i,j}$. One can easily check that $\xi_i(S_i) = 1$ for all i and $\xi_i(S_j) = 0$ when $i \neq j$. The measure of a set being the infimum of the measures of open sets covering it, for each i, there is an open set U_i covering S_i such that $\xi_j(U_i) \leq 2^{-s-1}\varepsilon$ for $i \neq j$ (and of course, $\xi_i(U_i) = 1$ for all i). Now, we use the fact that the measure an open set is the supremum of the measures of the clopen sets it contains. Therefore, there exists for each i a clopen set $U_i' \subseteq U_i$ such that $\xi_i(U_i') \geq 1 - \varepsilon/2$ (and of course $\xi_i(U_j') \leq 2^{-s-1}\varepsilon$ for $i \neq j$). Now, it suffices to define for each i: $V_i = U_i' \setminus \bigcup_{j \neq i} U_j'$, which is a clopen set for each i, of ξ_i-measure at least $1 - \varepsilon/2 - 2^s \cdot 2^{-s-1}\varepsilon = 1 - \varepsilon$. The pairwise disjointness of the V_i is clear from their definition.

We can therefore also compute the maximum of the granularities of all the clopen sets $V_j^{(i)}$ for $i \leq k(\varepsilon)$ and $j \leq 2^s$. Denote this maximum by $L(n)$ and finally define the function f by $f(n) = 2^{-L(n)-s(n)}$.

Suppose now that for every non-empty $\mathcal{B}' \subseteq \mathcal{B}$, there exists some $\mu \in \mathcal{B}'$ and some n,

$$\mu(\mathrm{PREC}(\mathfrak{A}, f, n)) > 1 - \eta - \delta/2$$

for some measure $\mu \in \mathcal{B}$ and some $n \geq N$. Set again $\varepsilon = \min(2^{-n} \cdot \delta/4, r)$ and compute a family $\mathcal{D}_1, \mathcal{D}_2, \cdots \mathcal{D}_{k(\varepsilon)}$ covering $\mathcal{M}_0$ so that for any ball $\mathcal{B}$ of radius at least ε, there is some $\mathcal{D}_i \subseteq \mathcal{B}$.

Recall that $\mathrm{PREC}(\mathfrak{A}, f, n)$ is a clopen set of granularity n. Thus, if $\rho(\nu, \mu) < 2^{-n} \cdot \delta/4$, then $\nu(\mathrm{PREC}(\mathfrak{A}, f, n)) > 1-\eta-\delta/2-\delta/4 = 1-\eta-3\delta/4$. And thus, by definition of the $\mathcal{D}_i$, there exists i such that for all $\nu \in \mathcal{D}_i$, $\nu(\mathrm{PREC}(\mathfrak{A}, f, n)) > 1-\eta-3\delta/4$. Moreover, such an i can be found effectively knowing $\mathrm{PREC}(\mathfrak{A}, f, n)$ and δ. Fix such an i and set $\mathcal{D} = \mathcal{D}_i$.

Now run the algorithm $\mathfrak{A}$ on all possible strings σ of length n. On some of these strings, the algorithm has precision less than $f(n)$; we ignore such strings. On some others, $\mathfrak{A}(\sigma)$ has precision at least $f(n)$ and thus returns a sequence containing some ball $\mathcal{A}$ of radius $< f(n)$. Call $\mathcal{A}_1, ..., \mathcal{A}_t$ all such balls (obtained by some $\mathfrak{A}(\sigma)$ with σ of length n). Note that $t \leq 2^n$. Let $\alpha_1, ..., \alpha_t$ be the centers of these balls, and consider their average $\beta = (1/t)\sum_{i \leq t} \alpha_i$. Since the V_i are disjoint, by the pigeonhole principle, there exists some j such that such that $\beta(V_j) \leq 2^{-s}$, and thus $\alpha_i(V_j) \leq t \cdot 2^{-s} \leq 2^{n-s}$ for all i. Fix such a j and set $V = V_j$, and $\xi = \xi_j$.

Recalling that the granularity of V is at most $L(n)$, we can apply the randomness deficiency lemma, we have for all $X \in V$ and all i:

$$\begin{aligned}
\mathbf{ed}(X|\mathcal{A}_i) &\geq -\log\left(\alpha_i(V) + 2^{L(n)} f(n)\right) - K(\alpha, V, f(n)) - O(1) && (1)\\
&\geq -\log\left(2^{n-s(n)} + 2^{L(n)} f(n)\right) - K(\alpha, V, f(n)) - O(1) && (2)\\
&\geq s(n) - n - K(\alpha, V, f(n)) - O(1) && (3)
\end{aligned}$$

It remains to evaluate the complexities of α_i, V and $f(n)$. The measure α_i is the center of the first ball of radius $< f(n)$ output by a computation $\mathfrak{A}(\sigma)$ for some σ of length n. Thus $K(\alpha_i, f(n)) \leq K(f(n))+n+O(1) \leq K(s(n))+n+O(1)$. To get a good upper bound for $K(V)$, we need to observe that $K(V)$ can be computed from $\mathrm{PREC}(\mathfrak{A}, f, n)$, which, being a clopen set of granularity at most n, has complexity at most $2^{n+O(1)}$. Indeed, knowing this set, one can compute the open set of measures ν such that $\nu(\mathrm{PREC}(\mathfrak{A}, f, n)) > 1-\eta-3\delta/4$ and effectively find a ball $\mathcal{D}$ as above. Then, from $\mathcal{D}$, the sequence of clopen sets $V_1, \ldots, V_{2^s}$ can be effectively computed. Moreover, to choose the V as above, we need to know β, hence the sequence of measure $\alpha_1, \ldots \alpha_t$. But these can also be found knowing $\mathrm{PREC}(\mathfrak{A}, f, n)$, by definition of the latter. Thus we have established that $K(V) \leq 2^{n+O(1)}$.

Plugging all these complexity estimates in (3), we get

$$\begin{aligned}
\mathbf{ed}(X|\mathcal{A}_i) &\geq s(n) - 2n - K(s(n)) - O(1) && (4)\\
&\geq s(n) - 2\log(s(n)) - 2n - O(1) && (5)
\end{aligned}$$

Thus, by taking $s(n) = 4n + d$ for some large enough constant d, we get that

$$\mathbf{ed}(X|\mathcal{A}_i) > n$$

for all $X \in V$ and all i. But the clopen set V has ξ-measure at least $1 - \delta/8$, so by definition of the $\mathcal{A}_i$, $\mathfrak{A}$ returns a ξ-inconsistent answer for deficiency level n on a set of ξ-measure at least $1 - \eta - 3\delta/4 - \delta/8$ of strings of length n. That is,

on a set of ξ-measure at least $1 - \eta - 3\delta/4 - \delta/8$, we there is a $\mathcal{B} \in \mathfrak{A}(X \upharpoonright n)$ such that $\mathbf{ed}(X|\mathcal{A}_i) > n$. Note that this is a Σ^0_1-property of ξ, so we can in fact effectively find a ball $\mathcal{B}''$ of measures on which this happens. For every $\nu \in \mathcal{B}''$, $\mathfrak{A}(\sigma)$ is null on a set of strings of ν-measure at least η (by assumption) and is inconsistent on a set of measure at least $1 - \eta - 7\delta/8$, so $\mathrm{SUCC}(\mathfrak{A}, n, n,)$ has a ν-measure of at most $7\delta/8$, which is what we wanted.

Now, we have reached our first goal which was to show that some $\mathcal{U}_m$ is not dense in $\mathcal{B}$ for some m. Let N' be such an m. Note that the $\mathcal{U}_m$ are non-increasing so this further means that there is a ball $\mathcal{B}' \subseteq \mathcal{B}$ which does not intersect any of the $\mathcal{U}_n$ for $n \geq N'$. By definition, this means that on any measure ν of that ball $\mathcal{B}'$, the algorithm does not reach precision $f(n)$ on a set of strings of ν-measure at least $\eta + \delta/2$. Thus, it suffices to consider the algorithm $\mathfrak{A}'$ which on any input σ does the following: it runs $\mathfrak{A}(\sigma)$ until $\mathfrak{A}(\sigma)$ reaches precision $f(|\sigma|)$. If this never happens, $\mathfrak{A}'(\sigma)$ remains null. If it does, then $\mathfrak{A}'(\sigma)$ returns the same list of balls as $\mathfrak{A}(\sigma)$. Clearly the algorithm $\mathfrak{A}'$ is δ-good since for every σ in the domain of $\mathfrak{A}$, $\mathfrak{A}'(\sigma) = \mathfrak{A}(\sigma)$. But by construction our new algorithm $\mathfrak{A}'$ is such that $\nu(\mathrm{NULL}(\mathfrak{A}', N')) \geq \eta + \delta/2$ for all $\nu \in \mathcal{B}'$. This finishes the proof. □

References

CV13. Nick Chater and Paul Vitanyi. Algorithmic identification of probabilities. http://arxiv.org/abs/1311.7385, 2013.

DH10. Rodney Downey and Denis Hirschfeldt. *Algorithmic randomness and complexity*. Theory and Applications of Computability. Springer, 2010.

Gác05. Peter Gács. Uniform test of algorithmic randomness over a general space. *Theoretical Computer Science*, 341(1-3):91–137, 2005.

LV08. Ming Li and Paul Vitányi. *An introduction to Kolmogorov complexity and its applications*. Texts in Computer Science. Springer-Verlag, New York, 3rd edition, 2008.

Nie09. André Nies. *Computability and randomness*. Oxford Logic Guides. Oxford University Press, 2009.

ZZ08. Thomas Zeugmann and Sandra Zilles. Learning recursive functions: a survey. *Theoretical Computer Science*, 397:4–56, 2008.